

CSci 5299

Senior Thesis

**Parallel Implementations of Neural Network Training:
Two Back-Propagation Approaches**

**Jeffrey Dean
ID #1321294**

Introduction

The back-propagation algorithm for training neural networks is used in a variety of applications due to its simplicity and ease of implementation. This paper assumes the reader is somewhat familiar with neural networks in general and the back-propagation algorithm in particular. An introduction to these topics can be found in [Rume 86]. The simplicity and widespread use of the back-propagation algorithm make it a logical starting point for developing a parallel algorithm for training neural networks. The work presented in this paper focuses on two back-propagation based approaches to training a neural network in parallel.

The first approach is a pattern partitioning approach. In this approach, the entire network is represented on each processor and the set of input patterns is divided among the available processors. Each processor independently computes the weight changes necessary to minimize the error for its subset of the input patterns. After each processor has computed these weight changes, they are combined and broadcast to all processors.

The second approach involves partitioning the neurons of the neural network among the available processors and having each processor propagate the entire set of input patterns. The advantage of this approach is that changes do not need to be combined among the processors. However, communication is still needed during the propagation phase since each processor now represents only a portion of the entire network.

The pattern partitioning approach was implemented first due to its simplicity. Given n input patterns and p processors, each processor was given n/p input patterns (with appropriate rounding for non-integer values, of course) and the entire set of weights for the network. During each epoch, each processor propagates its input patterns through the network, and then back-propagates the error values. At some point, the processors need to combine the delta values resulting from the back-propagation phase.

Two different approaches to this combination were attempted. The first involved saving the results of the pattern propagation and the delta values from the back-propagation in matrices and then combining these matrices. This turned out to be a naive approach since the amount of communication that needed to be done increased linearly with the number of input patterns. A better approach was to collect the weight delta values for all of a processor's n/p patterns and then to combine these weight delta matrices by summing them. This approach provided much better scalability.

The network-partitioned approach that was implemented involves dividing the network elements among the p processors. The model of a software pipeline is used, with

the processors communicating in a ring fashion. Patterns are placed in one end of the pipeline, and as they flow through the pipeline are processed by each processor's neurons. One processor does no actual processing of neurons but acts as a monitoring processor, placing messages into the pipeline, retrieving them from the end of the pipeline and keeping track of bookeeping tasks. The network elements are divided among the other $p-1$ processors.

Measures of Efficiency

The performance of a parallel algorithm can be split into two components: the time spent doing useful work, and the time spent idling or in communications. Call these two components t_{calc}^i , the time spent doing useful calculations, and t_{comm}^i , the time spent in communications and not performing useful work (the superscript i denotes the processor number). The sum of t_{calc}^i for all processors is defined as T_{calc} , and the sum of t_{comm}^i for all processors is T_{comm} . The total time taken by an algorithm on p processors is:

$$T = \sum_{i=1}^p (t_{calc}^i + t_{comm}^i) = T_{calc} + T_{comm} = pT_p$$

where T_p is the execution time on p processors.

Problem Size W : The problem size is defined as the total amount of computation done by the best known sequential algorithm.

Speedup S : The speedup on p processors is defined as ratio of $\frac{T_1}{T_p}$, where T_1 is the time of the best sequential algorithm.

Efficiency E : Efficiency measures the effective utilization of p processors and is defined as $\frac{S_p}{p}$. Essentially measures what percentage of a linear speedup was obtained on p processors.

Iso-efficiency - The iso-efficiency function describes the relationship between W and p that must be maintained in order to keep a constant efficiency. See [Gupt 1990] for a complete discussion.

Definitions

In order to analyze the efficiency of these algorithms, we need some common definitions.

The inputs to a neural network consist of:

n input patterns, and

m neuron layers, each of size k , along with $(m - 1)$ weight matrices of size $k \times k$ connecting the layers. Note that we are assuming that each layer has the same number of neurons for the purposes of our analysis. The programs that were developed allow differing number of neurons for each layer. This is generally the case in actual neural network applications.

We are given p processors.

From these values, we can derive some useful expressions:

M , the sizes of all weight matrices, is $(m - 1)k^2$.

C , the number of hidden and output layer neurons, is $(m - 1)k$.

Pattern-Partitioned Approach

The pattern partitioned approach on a hypercube is fairly simple to implement. The host processor sends the weight matrices for the entire network to each processor, along with the subset of n/p input patterns that each processor will be handling. During an epoch, each processor calculates the weight delta values for its pattern set by propagating its inputs through the network. These weight delta matrices are then sent up a tree (mapped onto a hypercube in the usual manner). At each non-leaf node, the two sets of matrices are combined using matrix addition. The node at the root of the tree then has the weight delta matrices for the entire input pattern set, and broadcasts these to all the other processors (in log time, using the typical hypercube broadcasting method). Each processor then updates its weights and begins processing its input patterns for the next epoch.

The sequential algorithm takes time proportional to M for each input pattern. Thus, the sequential algorithm's time per epoch is $O(nM)$.

Analysis

The computation time per processor is:

$$t_{calc}^i = \frac{nM}{p}$$

The communication time per processor is:

$$t_{comm}^i = M \log p$$

Thus, the speedup of the pattern-partitioned approach is:

$$S = \frac{nM}{\frac{nM}{p} + M \log p}$$

The efficiency is:

$$E = \frac{nM}{nM + M p \log p} = \frac{1}{1 + \frac{p \log p}{n}}$$

In order to maintain constant efficiency, $\frac{p \log p}{n}$ must remain constant. Thus, the iso-efficiency function is:

$$n \propto p \log p$$

Experimental Results

Several different sets of input data were generated in order to evaluate the performance of the pattern-partitioned approach. The intuitive result that better speedups would be obtained with larger numbers of input patterns held. Speedup curves for three of these input sets are depicted below:

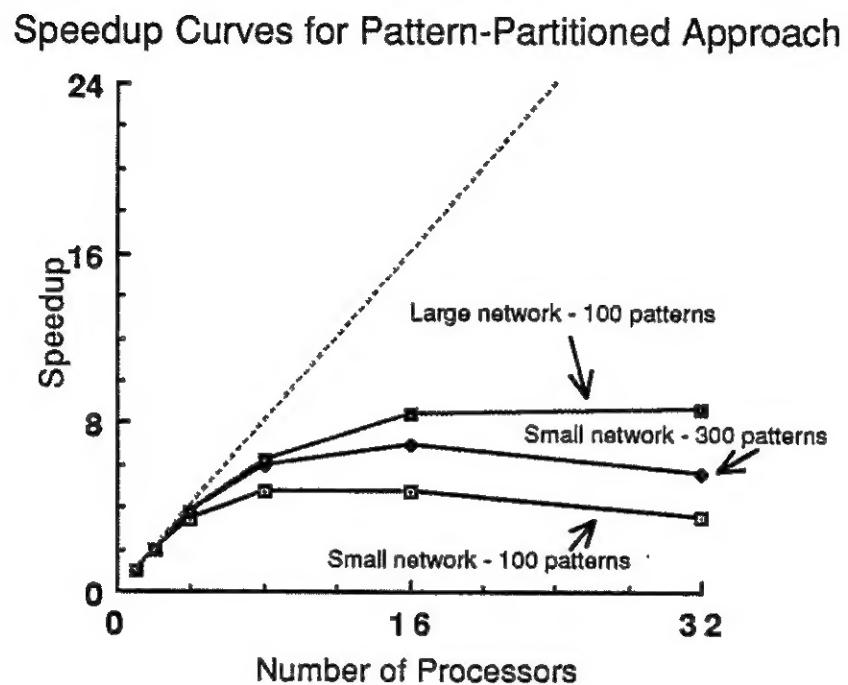


Fig. 1. Speedup curves for pattern-partitioned approach. The small network contained 3 layers with 9, 4 and 2 neurons. The large network contained 3 layers with 10, 21 and 10 neurons.

Pipelined Approach

The second method that was implemented involved a network-partitioning, rather than a pattern-partitioning approach approach. During an epoch, a given pattern will be handled by every processor in the system, but that processor will only do a portion of the processing for that pattern. Only hidden and output layer nodes are represented on processors; since the input layer has no weights incident on it, any processing of input layer nodes is assumed to be handled by the first hidden layer. One processor is assigned the role of manager and is responsible for filling the pipeline with messages, collecting the output at the other end of the pipeline, and various bookeeping tasks. The C hidden and output layer neurons are evenly divided among the remaining $p - 1$ processors¹. Initially, the managing processor generates messages to completely fill the pipeline; thereafter, it receives output from the end of the pipeline and continues to ensure that the pipeline is full. For a complete discussion of the algorithm, see [Alle 89].

Each message in this approach contains the input pattern, target output pattern, a placeholder for the input to and output from the current layer, and a set of m vectors to hold the error information for each layer. The size of each message is $O(mk)$.

Analysis

The computation time per processor per epoch is:

$$t_{calc}^i = \frac{n}{p}M = \frac{n}{p}mk^2$$

The communication time per processor epoch is:

$$t_{comm}^i = nC = nmk$$

Thus, the speedup of the pipelined network-partitioned approach is:

$$S = \frac{nmk^2}{\frac{nmk^2}{p} + nmk}$$

¹Allen and Saha [ALLE 89] conducted experiments which showed that output layer neurons take 3 times as long to process as hidden layer neurons, so some load balancing may be necessary to offset this difference. However, this was not taken into account in this implementation.

The efficiency is:

$$E = \frac{nmk^2}{nmk^2 + pnmk} = \frac{1}{1 + \frac{p}{k}}$$

In order to maintain constant efficiency, $\frac{p}{k}$ must remain constant. Thus, the iso-efficiency function is:

$$k \propto p$$

Experimental Results

Several neural networks of various sizes were trained using this pipelined approach. Due to a limitation on message sizes between processors on the NCUBE multicomputer, results were only analyzed up to 32 processors. The calculations do not include the managing processor, since its functions could also be handled by the host processor. Thus, on a 3-dimensional hypercube, 7 processors were performing training, so the speedup is presented for 7 processors, not 8 processors. On the graph below, there are two curves: one is for the pipelined approach of a moderate sized network and the other is for the same inputs under the pattern-partitioned approach described above. Speedup curves for both methods are shown.

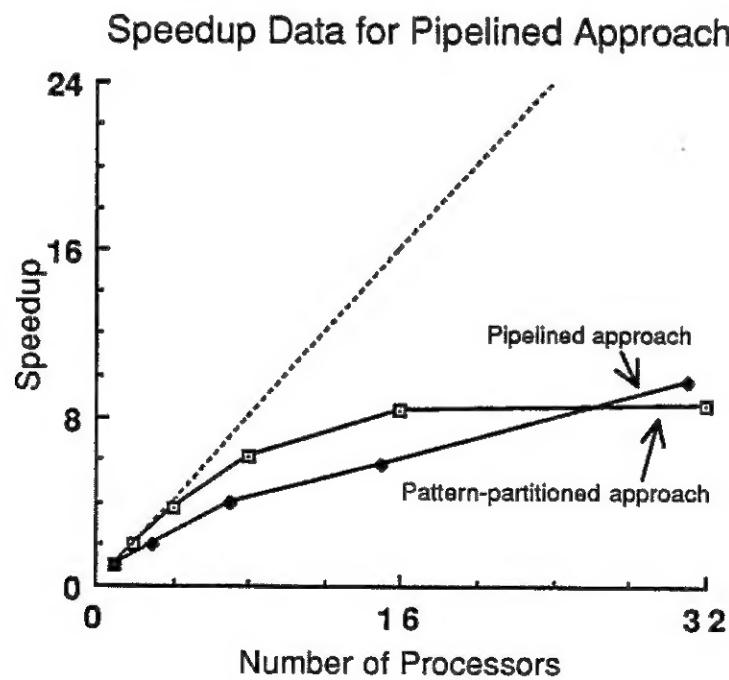


Fig 2. Data for large network dataset (3 layers, 10, 21 and 10 neurons per layer) for pipelined and pattern-partitioned approaches.

Discussion

Both approaches seemed to perform fairly well. The pattern-partitioned approach is less scalable than the pipelined approach, but performs better with a smaller number of processors. The pipelined approach could be improved in several ways. One way would be to balance the processing times taken for each processor in the pipeline by performing some analysis on the time taken for hidden and output layer neurons rather than simply dividing the available neurons evenly among the processors. This would eliminate any bottlenecks caused by differences in processor speed, and would increase concurrency. Another possibility is to have a two-way pipeline, in which messages flow to the end of the network, and then flow backwards to implement the back-propagation phase rather than saving the results of the previous epoch to be cycled through in the next epoch. This would reduce the message size by a factor of m and would undoubtedly increase the efficiency of the algorithm.

References

- [Alle 89] Allen, Wayne and Avijit Saha, *Parallel Neural-Network Simulation Using Back-propagation for the ES-Kit Environment*, Microelectronics and Computer Technology Corporation, Austin, Texas, 1989.
- [Ghos 89] Ghosh, Joydeep, and Kai Hwang, "Mapping Neural Networks onto Message-Passing Multicomputers," *Journal of Parallel and Distributed Computing*, April 1989.
- [Gupt 90] Gupta, Anshul and Vipin Kumar, *On the Scalability of FFT on Parallel Computers*, Dept. of Computer Science, University of Minnesota, 1990.
- [Pome 88] Pomerleau, Dean A., George L. Gusciora, David S. Touretzky and H.T. Kung, "Neural Network Simulation at Warp Speed: How We Got 17 Million Connections per Second," Internation Conference on Systolic Arrays, 1988.
- [Rume 86] Rumelhart, D. E., and J.L. McClelland, editors, "Parallel Distributed Processing: Explorations in the Microstructure of Cognition," Bradford Books/MIT Press, Cambridge, Mass., 1986.

12.16.02

09/06/06

net.h

```
/* net.h
 */
/* #define DEBUG */
/* #define QUICKF */
/* #define IBM */

#define ESC 27
#define ERROLEVEL 0.02
#define ITEMS 8
#define MAXLAYERS 5
#define FALSE 0
#define TRUE 1
#define MAXDATA 15000
#define MAXBUFSIZE ((MAXDATA+4)*sizeof(float))
#define MAXCUBDIM 6
#define MS_PER_TICK 0.166
#define LBUFFLEN 512
#define SBUFFLEN 10

/* message types for inter-processor communication */
#define M_INITIAL 1
#define M_PATTERN 2
#define M_WEIGHTS 3
#define M_UPDATE 4
#define M_OUTPUT 5
#define M_ERRORS 6
#define M_ERRORS2 7
#define M_TIME 8

/* typedefs and prototypes for dynamic storage of arrays */
typedef float *PFLOAT;
typedef PFLOAT VECTOR;
typedef PFLOAT MATRIX;
typedef struct layerrec {
    int lsize;
    int wsize;
    VECTOR del;
    MATRIX wdel;
    MATRIX wi;
    VECTOR o;
} LAYER;

typedef struct vectormsg {
    int v_id;
    int v_cols;
    float v_data[MAXDATA];
} VECTORMSG, *VMSG_PTR;

typedef struct matrixmsg {
    int m_id;
    int m_rows;
    int m_cols;
    float m_data[MAXDATA];
} MATRIXMSG, *MM_PTR;

typedef struct initialmsg {
    int i_id;
    int i_layers;
    int i_iters;
    int i_npats;
    int i_startat;
    int i_endpat;
    long i_seed;
    float i_eta;
    float i_alpha;
    float i_epsilon;
    int i_lsize[MA
```

network.c

12:16:26

```

/* network.c

Common routines to both the host and node programs. Contains
routines for allocating vectors and matrices, and for sending
vectors and matrices as messages.

*/
#include <stdio.h>
#define NULL 0
Definitions for stream input/output.

Copyright (c) Borland International 1987,1988
All Rights Reserved.

/*
 * include <math.h>
 * include <cctype.h>
 * include <stdlib.h>
 * include <conio.h>
 * include "math.h"
 * include "cctype.h"
 * include "string.h"
 * include "net.h"
 */

----- Vector and Matrix message passing routines -----
----- IBM -----


#ifndef IBM
int SendVector(VECTOR vector, int nCols, VM_PTR mbuf, int dest, int msg_id)
{
    int nRows;
    VECTOR vector;
    int nCols;
    VM_PTR mbuf;
    int dest;
    int msg_id;
}

register int i;
int buflen, flag;
long result;

mbuf->v_cols = nCols;
mbuf->v_id = msg_id;
#endif IBM
swap (mbuf->v_cols, sizeof (int));
swap (mbuf->v_id, sizeof (int));
for (i = 0; i < nCols; i++)
{
    mbuf->v_data[i] = vector [i];
}
#endif IBM
swap (mbuf->v_data[i], sizeof (float));
buflen = 2 * sizeof (int) + nCols * sizeof (float);
#endif IBM
result = nwrite (chan, (char *)mbuf, buflen, dest, msg_id);
return (result);
else
nwrite ((char *)mbuf, buflen, dest, msg_id, &flag);
return (flag);
}

```

```

swap (&mbuf->m_id, sizeof (int));
#endif
for (i = 0; i < nRows; i++)
    for (j = 0; j < nCols; j++)
        {
            mbuf->m_data[i*nCols+j] = matrix [i][j];
        }
#endif HOST
swap (& (mbuf->m_data[1*nCols+1]), sizeof (float));
#endif
buflen = 3 * sizeof (int) + nRows * nCols * sizeof (float);
}

ifdef HOST
    result = fwrite (chan, (char *)mbuf, buflen, dest, msg_id);
else
    fwrite ((char *)mbuf, buflen, dest, msg_id, &flag);
endif
return (flag);
}

ifdef HOST
    result = fread (chan, (char *)mbuf, buflen, dest, msg_id);
else
    fread ((char *)mbuf, buflen, dest, msg_id, &flag);
endif
}

ifdef IBM
int GetMatrix(MATRIX matrix, int nRows, int nCols,
              int from, int msg_id)
{
    int GetMatrix(MATRIX matrix, nRows, nCols, mbuf,
                  from, msg_id)
MATRIX matrix;
int nRows;
int nCols;
MM_PTR mbuf;
int from;
int msg_id;
}
#endif
register int i, j;
int buflen, msgnode, type, flag;
long result;

buflen = 3 * sizeof (int) + nRows * nCols * sizeof (float);
msgnode = from;
type = msg_id;
#endif HOST
result = fread (chan, (char *)mbuf, buflen, msgnode, &type);
else
    fread ((char *)mbuf, buflen, msgnode, type, &flag);
for (i = 0; i < nRows; i++)
    for (j = 0; j < nCols; j++)
        {
            ifdef HOST
                swap (& (mbuf->m_data[1*nCols+1]), sizeof (float));
            endif
            matrix [i][j] = mbuf->m_data[i*nCols+j];
        }
#endif HOST
return (result);
else
    return (flag);
}

/*----- Array storage allocation routines -----*/
/*----- Allocate space for vector of float cells for
   -----*/

```

90/06/06
12:16:26

network.C

```
free(matrix[1]);
free(matrix);
}

#ifndef IBM
void LayerAllocate(LAYER *layer, int layersize, int wtsize)
{
    void LayerAllocate(layer, layersize, wtsize)
    LAYER *layer;
    int layersize;
    int wtsize;
    endiff
    {
        layer->size = layersize;
        layer->wsize = wtsize;
        /* w and wdel must contain room for the node bias */
        MatrixAllocate (&(layer->w), layersize, wtsize+1);
        MatrixAllocate (&(layer->wdel), layersize, wtsize+1);
        VectorAllocate (&(layer->o), layersize);
        VectorAllocate (&(layer->del), layersize);
    }

    #ifdef IBM
    void LayerFree(LAYER *layer)
    {
        LayerFree(layer)
        LAYER *layer;
        endiff
        {
            MatrixFree (layer->w, layer->size);
            MatrixFree (layer->wdel, layer->size);
            free (layer->o);
            free (layer->del);
        }
    }

    #ifdef QUICKF
    static int doneinit = FALSE;
    static VECTOR v;
    static int vsize = 1000;
    static float lo = -5.0;
    static float hi = 5.0;
    static float step;
    register int i;
    float Y, Y2, d;
    endiff
    if (!doneinit)
    {
        VectorAllocate (&v, vsize);
        step = (hi - lo) / vsize;
        for (i = 0, Y = lo; i < vsize; i++, Y+=step)
            v[i] = 1.0 / (1.0 + exp (-Y));
        doneinit = TRUE;
    }
    if (x < lo || x > hi)
        return (1.0 / (1.0 + exp (-x)));
    else

```

3

90:06/06
12:16:16

nethost.c

```
#include <stdio.h>
#include <stdlib.h>
#include <conio.h>
#include <math.h>
#include <cctype.h>
#include <string.h>
#include "net.h"

def IBM
    ShowMatrix(MATRIX m, int nRows, int nCols);
    se
        ShowMatrix();
    dif

define storage for net layers */
Arrays for inputs, outputs, deltas, weights & targets */
ER    net(MAXLAYERS);
Lsize[MAXLAYERS];
TOR PatternID;
RIX target; /* target output */
RIX input; /* input patterns */
at   b[SMALLBUFSIZE];
at   msSize;
at   avgTime;
at   argc;
at   argv[];

Cube communication variables */
chan;
cubedim, cubesize;
types;
*mbuf, *lbuf, *nflags;
b [SMALLBUFSIZE];
szResults;
szWeights;
anPatterns,
nIterations,
eta,
alpha,
nReportErrors = 1;
float ErrorLevel = ERRORLEVEL;
float ErrorLevel;
float pat_node;
float pat_num;
int i;
int j;
int l;
int n;

#include "network.c"

def IBM
    d main(int argc, char *argv[])
    se
        d main(argc, argv)
            argc;
            argv[];
        dif

        float eta = 0.15; /* default learning rate
                           /* default momentum factor
                           */
        float alpha = 0.075; /* error reporting frequency
                           /* error level */
        int nReportErrors = 1; /* satisfactory error value
                           */
        float ErrorLevel; /* latest sum squared error value
                           */
        float pat_node; /* patterns per node
                           */
        float pat_num; /* current pattern number
                           */
        int i; /* index variable
                           */
        int j; /* layer index
                           */
        int l; /* neuron index
                           */

        /* get sizes of layers */
        for (l = 0; l < nLayers; l++)
            fscanf(fpInp, "%d", &Lsize[l]);
        printf("%d layers, %d patterns, %d layers, %d patterns);\n"
               "index pattern number\n"
               "patterns, /* number of patterns desired */
               "nLayers, /* number of layers in network */
               "nIterations; /* number of iterations desired */
               FILE *fpInp,
               *fpResults,
               *fpWeights;
               szResults[66];
               char szWeights[66];
               char *programme = *argv; /* name of executable DOS 3.x only */
               char c;

               /* allocate message buffer */
               mbuf = (char *)malloc(MAXBUFSIZE);

               printf ("Neural net simulation\n");
               /* read optional - arguments */
               for (l; argc > 1; argc--)
                   char *arg = *++argv;
                   if (*arg != '-')
                       break;
                   switch (*++arg)
                   {
                       case 'e': sscanf(*+arg, "%f", &nErrorLevel); /*ErrorLevel*/;
                       case 'd': sscanf(*+arg, "%f", &nEta); /*Eta*/;
                       default: break;
                   }
               if (argc < 2)
                   {
                       fprintf(stderr, "Usage: %s (-en -df) runfilename\n", programme);
                       fprintf(stderr, " -en => report error every n iterations\n");
                       fprintf(stderr, " -df => done if mean squared error < f\n");
                       exit(1);
                   }
               /* Open run file for reading */
               if (fpInp = fopen(*argv, "r")) == NULL)
                   {
                       fprintf(stderr, "%s: can't open file %s\n", programme, *argv);
                       exit(1);
                   }
               /* read and parse the run specification line */
               fscanf(fpInp,
                      "%s %d %d %f %d",
                      szResults, /* output results file */
                      szWeights, /* initial weights file */
                      anPatterns, /* number of patterns to learn */
                      nIterations, /* number of iterations through the data */
                      eta, /* learning rate */
                      alpha, /* momentum factor */
                      nLayers); /* number of layers in network */

               /* get sizes of layers */
               for (l = 0; l < nLayers; l++)
                   fscanf(fpInp, "%d", &Lsize[l]);
               printf("%d layers, %d patterns, %d layers, %d patterns);\n"
                      "index pattern number\n"
                      "patterns, /* number of patterns desired */
                      "nLayers, /* number of layers in network */
                      "nIterations; /* number of iterations desired */
                      FILE *fpInp,
                      *fpResults,
                      *fpWeights;
                      szResults[66];
                      char szWeights[66];
                      char *programme = *argv; /* name of executable DOS 3.x only */
                      char c;
```

90/06/06
12:16:16

nefhost.c

```

for (l = 0; l < nLayers; l++)
    printf ("%d", lSize [l]);
    printf ("\n");

/*-----allocate dynamic storage for all data -----*/
for (l = 0; l < nLayers; l++)
    LayerAllocate (&net [l], lSize [l], (l==0) ? 0:lSize [l-1]);
    MatrixAllocate (starget,
                    MatrixAllocate (&input,
                                  nPatterns,
                                  lSize [nLayers-1]);
                    VectorAllocate (&PatternID,
                                  nPatterns,
                                  nPatterns);
    /*----- Read the initial weight matrices: -----*/
    if ((fpWeights = fopen (szWeights, "r")) == NULL)
    {
        fprintf (stderr, "%s: can't open file %s\n",
                progname, szWeights);
        exit (1);
    }

    /* read weights */
    for (l = 0; l < nLayers; l++)
        for (n = 0; n < net [l].lSize; n++)
            for (i = 0; i < net [l-1].lSize; i++)
            {
                fscanf (fpWeights, "%f",
                        &net [l].w[n][i]);
                net [l].wdei[n][i] = 0.0;
            }
    fclose (fpWeights);

    /*----- Read in all patterns to be learned -----*/
    for (p = 0; p < nPatterns; p++)
    {
        for (i = 0; i < net [0].lSize; i++)
            if (fscanf (fpInp, "%f",
                       &input [p][i]) != 1)
                goto ALPATTERNSREAD;

        /* read in target outputs for input patterns read */
        for (i = 0; i < net [nLayers-1].lSize; i++)
            fscanf (fpInp, "%f",
                    &targer [p][i]);
        fscanf (fpInp, "%f",
                &PatternID[p]);
    }

    LPATTERNSREAD:
    if (p < nPatterns)
    {
        fprintf (stderr, "%s: Id out of %d patterns read\n",
                progname, p, nPatterns);
        nPatterns = p;
    }
}

cubedim = MAXCUBDIM + 1;
while ((cubedim < 0) || (cubedim > MAXCURDIM))
{
    printf ("Enter Cube Dimension {0 - %d} : ", MAXCURDIM);
    scanf ("%d", &cubedim);
}

cubedim = power (2, cubedim); /* number of processors */
lbuf = (char *)malloc (IBUFLEN);
printf ("Allocating hypercube.\ntry ");
i = 1;

cubedim = MAXCUBDIM + 1;
while ((cubedim < 0) || (cubedim > MAXCURDIM))
{
    printf ("%d", lbuf, nFlags, -cubesize, lbufLen);
    free (lbuf);
}

printf ("Rode program loaded to %d processors.\n", cubesize);
printf ("Setting up Initial message to nodes.\n");
printf ("Set up Initial message. Now swapping.\n");
swap (&((IM_PTR)mbuf)->l_layers, sizeof (int));
swap (&((IM_PTR)mbuf)->l_iters, sizeof (int));
swap (&((IM_PTR)mbuf)->l_inputs, sizeof (int));
swap (&((IM_PTR)mbuf)->l_seeds, sizeof (int));
swap (&((IM_PTR)mbuf)->l_eta, sizeof (float));
swap (&((IM_PTR)mbuf)->l_alpha, sizeof (float));
swap (&((IM_PTR)mbuf)->l_epsilon, sizeof (float));
printf ("Set up Initial message. Now swapping.\n");
swap (&((IM_PTR)mbuf)->l_stacks, sizeof (float));
swap (&((IM_PTR)mbuf)->l_patterns, sizeof (int));
swap (&((IM_PTR)mbuf)->l_seeds, sizeof (long));
swap (&((IM_PTR)mbuf)->l_alpha, sizeof (float));
swap (&((IM_PTR)mbuf)->l_epsilon, sizeof (float));
for (l = 0; l < MAXLAYERS; l++)
{
    ((IM_PTR)mbuf)->l_size [l] = lSize [l];
    swap (&((IM_PTR)mbuf)->l_lsize [l], sizeof (int));
}
pat_node = (float)lPatterns / cubesize;
pat_num = 0;
printf ("Now sending initial message to nodes. Total patterns = %d\n",
        nPatterns);

for (l = 0; l < cubesize; l++)
{
    /* Synchronization step */
    type = M_INITIAL;
    nread (chan, (char *)b, sizeof (int), &l, &type);
    ((IM_PTR)mbuf)->l_startpat = (int) pat_num;
    ((IM_PTR)mbuf)->l_endpat = (int) (pat_num + pat_node - 1);
    printf ("%d gets patterns %d to %d.\n", l,
            ((IM_PTR)mbuf)->l_startpat,
            ((IM_PTR)mbuf)->l_endpat);
    swap (&((IM_PTR)mbuf)->l_endpat, sizeof (int));
    pat_num += pat_node;
}

type = M_INITIAL;
nwrite (chan, (char *)mbuf, sizeof (INITIALMSG), 1, type);
printf ("Initial message sent. Sending Input matrix.\n");
printf ("Node %d", i);
printf ("Input, nPatterns, lsize [0], (IM_PTR)mbuf,
        1, M_PATTERN);
for (l = 0; l < cubesize; l++)
{
    printf ("%d..", l);
    SendMatrix (input, nPatterns, lsize [0], (IM_PTR)mbuf,
                1, M_PATTERN);
}

```

nethost.c

```

SendMatrix (target, nPatterns, lsize [nlayers-1], {MM_PTR} mbuf,
           i, M_OUTPUT);
for (j = 1; j < nlayers; j++)
  SendMatrix (net [j].W, lsize [j], lsize [j-1]+1, {MM_PTR} mbuf,
              i, M_WEIGHTS);

printf ("\n");
fprintf (stderr, nIterations > 1 ? "Training... \n" : "Testing\n");

/* Now we await the results */
do
{
  i = -1;
  type = M_UPDATE;
  nread (chan, (char *)bf, 4 * sizeof (float), &i, &type);
  for (j = 0; j < 4; j++)
    swap (bf+j, sizeof (float));
  printf ("Node %d: Iteration %5.0f/%-5.0f Error: %f, type = %5.0f\n",
          i, bf[0], bf[1], bf[2], bf[3]);
} while (bf [3] != 1);

/* Now get the weight matrices from node 0 */
for (l = 1; l < nlayers; l++)
  GetMatrix (net [l].W, lsize [l], lsize [l-1]+1, {MM_PTR} mbuf,
             0, M_WEIGHTS);

/* Now get the time taken from each node */
for (i = 0, avgTime = 0.0; i < cubesize; i++)
{
  type = M_TIME;
  nread (chan, (char *)b, sizeof (int), &i, &type);
  swap (b, sizeof (int));
  printf ("Node %d took %f ms (%d x %f) \n", i, b[0]*MS_PER_TICK,
          b[0], MS_PER_TICK);
  avgTime += b[0] * MS_PER_TICK;
}

printf ("\nAverage time per node: %6.1f ms\n", avgTime / cubesize);

printf ("Program terminating.\n");

/*----- free dynamic storage for data -----*/
for (l = 0; l < nlayers; l++)
  LayerFree (net [l]);
MatrixFree (target, nPatterns);
MatrixFree (input, nPatterns);
free (PatternID);
fclose (Pfilep);

/* close run file */

/*----- IBM -----*/
ifdef IBM
int ShowMatrix (MATRIX m, int nrows, int ncols)
else
int ShowMatrix (m, nrows, ncols)
MATRIX m;
int ncols;
int nrows;
endif
{
  int i, j;

  for (i = 0; i < nrows; i++)
    for (j = 0; j < ncols; j++)

```

90/06/06
12:16:18

netnode.c

```
/* netnode.c

Jeffrey Dean
ID 1321294
Term Project

Back-propagation neural network

*/
#define NODE

#ifndef IBM
#include <stdlib.h>
#endif
#ifndef MATH
#include <math.h>
#endif
#ifndef STRING
#include <string.h>
#endif
#ifndef NET_H
#include "net.h"
#endif

/* define storage for net layers */
LAYER net[MAXLAYERS];
int lsize[MAXLAYERS];
VECTOR patternID;
TARGET target;
MATRIX targetOutput;
MATRIX input;
MATRIX tempwdel;

long ransed = 1235734L;

ranseed = 1235734L;

/* Cube communication variables */
int type;
char *mbuf;
int b [SMALLBUFLEN];
bf [SMALLBUFLEN];

#include "network.c"

void main()
{
    float eta = 0.15, /* default learning rate
                        /* default momentum factor
    int nReportErrors = 1, /* error reporting frequency
    float ErrorLevel = ERRORLEVEL; /* satisfactory error level
    char MonitorError = 0; /* true when monitor error display
    float error;
    int i,
    j;
    int mask;
    int stepstowait;
    int msgnode;
    int l,
    n,
    p,
    q,
    npatterns,
    startPat,
    endPat,
    nInputNodes,
    nHiddenNodes,
    nOutputNodes,
    nLayers,
    maxLSize,
    nIterations;
    int time;
    whami (&nodeinfo, &proc, &host, &cnbedim);
    /* allocate message buffer */
    mbuf = (char *)malloc(MAXBUFSIZE);
    b[0] = 1;
    type = M_INITIAL;
    nwrite ((char *)b, sizeof (int), host, type, &flag);
    nread ((char *)mbuf, sizeof (INITIALMSG), host, type, &flag);
    nlayers = ((IM_PTR)mbuf)->layers;
    niterations = ((IM_PTR)mbuf)->literations;
    npatterns = ((IM_PTR)mbuf)->lnpatterns;
    startPat = ((IM_PTR)mbuf)->l_startpat;
    endPat = ((IM_PTR)mbuf)->l_endpat;
    ranseed = ((IM_PTR)mbuf)->l_seed;
    eta = ((IM_PTR)mbuf)->l_eta;
    alpha = ((IM_PTR)mbuf)->l_alpha;
    Errorlevel = ((IM_PTR)mbuf)->l_epsilon;
    maxSize = 0;
    for (i = 0; i < MAXLAYERS; i++)
    {
        lsize[i] = ((IM_PTR)mbuf)->lsize[i];
        maxSize = (lsize[i] > maxSize) ? lsize[i] : maxSize;
    }
    /* allocate dynamic storage for all data */
    for (i = 0; i < nlayer; i++)
    LayerAllocate(&net[i], lsize[i], (i==0) ? 0:lsize[i-1]);
    MatrixAllocate(&target,
    MatrixAllocate(&input,
    MatrixAllocate(&tempwdel,
    VectorAllocate(&patternID, npatterns),
    GetMatrix (input, npatterns, lsize[0], (IM_PTR)mbuf, host,
    M_PATTERN);
    GetMatrix (target, npatterns, lsize[nLayers-1], (IM_PTR)mbuf, host,
    M_OUTPUT);
    for (l = 1; l < nlayer; i++)
    {
        GetMatrix (net[i].M, lsize[i], lsize[0], (IM_PTR)mbuf, host,
        M_WEIGHTS);
        /* initialize weight deltas to 0 */
        for (n = 0; n < lsize[i]; n++)
        for (j = 0; j < lsize[i-1]; j++)
        net[i].wdel[n][j] = 0;
    }
    /* we now have all the matrices we need; start doing useful work */
    stime = ntime; /* get start time */
    /* begin iteration loop */
    for (q = 0; q < nIterations; q++)
    {
        /* initialize weight deltas to momentum factor */
        /* for (l = 1; l < nlayer; l++)
        for (n = 0; n < lsize[l]; n++)
        for (j = 0; j < lsize[l-1]; j++)
        */
        /* starting pattern number for node */
        /* ending pattern number for node */
        /* number of input nodes */
        /* number of hidden nodes */
        /* number of output nodes */
        /* number of layers in network */
        /* maximum layer size */
        /* number of iterations desired */
        /* Starting time */
    }
}
```

```

    net [l].wdel [n][j] = alpha * net[l].wdel [n][j];
}

/* initialize error level to 0 */
error = 0.0;

for (p = startPat; p <= endPat; p++)
{
    /*----- copy current pattern to outputs for level 0 -----*/
    for (n = 0; n < lsize [0]; n++)
        net [0].o[n] = input [p][n];
}

/*----- evaluate network using current weights -----*/
/* Sum outputs from layer to layer over all possible edges */
for (l = 1; l < nLayers; l++)
{
    for (n = 0; n < net[l].lsize; n++)
    {
        float sum = net[l].w[n][net[l-1].lsize];
        for (i = 0; i < net[l-1].lsize; i++)
            sum += net[l].w[n][i] * net[l-1].o[i];
        net[l].o[n] = f (sum);
    }
}

/* accumulate error term */
for (n = 0; n < lsize [nLayers-1]; n++)
{
    float temp = target [p][n] - net [nLayers-1].o [n];
    error += temp * temp;
}

/*----- calculate output deltas -----*/
/* Compute deltas for each output unit for the given pattern */
for (n = 0; n < net [nLayers-1].lsize; n++)
{
    float opn = net [nLayers-1].o [n];
    net [nLayers-1].delt[n] =
        (target [p][n] - opn) * opn * (1.0 - opn);
}

/*----- calculate hidden deltas -----*/
for (l = nLayers-2; l > 0; l--)
{
    for (n = 0; n < net [l].lsize; n++)
    {
        float sum = 0.0;
        float opn = net [l].o [n];
        for (i = 0; i < net [l+1].lsize; i++)
            sum += net [l+1].del [i] * net [l+1].w [i][n];
        net [l].del [n] = sum * opn * (1.0 - opn);
    }
}

/*----- adapt weights -----*/
for (l = nLayers - 1; l > 0; l--)
{
    for (n = 0; n < net [l].lsize; n++)
    {
        float dw;
        dw = eta * net [l].del [n] * net [l-1].o [n];
        net [l].wdel [n][l] += dw;
    }
}

/*----- calculate new weights -----*/
for (l = 0; l < net [l-1].lsize; l++)
{
    dw = eta * net [l].del [n] * net [l-1].o [n];
    net [l].wdel [n][l] += dw;
}

/*----- end of pattern loop -----*/
}

/*----- communicate deltas -----*/
/* first go up a binary tree (mapped onto the cube), sending */
/* the partial delta results for each node patterns. */
for (l = cubedim-1; l >= 0; l--)

mask = (l << 1);
if (nodeNo >= (mask << 1)) goto loopexit;
msgNode = nodeNo ^ mask;
if ((nodeNo & mask) == 0)
{
    /* accumulate error values from lower level nodes */
    type = M_ERRORS;
    read ((char *)bf, sizeof (float), msgNode, type, &flag);
    error += bf [0];
}

/* Get weight delta matrices for each layer and add to */
/* current weight deltas */
for (l = 1; l < nLayers; l++)
{
    GetMatrix (tempWdel, lsize [l], lsize [l-1]+1,
        (NM_PTR)mbuf, nodeNo ^ mask, M_WEIGHTS);
    for (n = 0; n < lsize [l]; n++)
        for (j = 0; j < lsize [l-1]; j++)
            net [l].wdel [n][j] += tempWdel [n][j];
}
else
{
    /* accumulate error values from lower level nodes */
    type = M_ERRORS;
    bf [0] = error;
    nwrite ((char *)bf, sizeof (float), msgNode, type, &flag);
    /* Send weight delta matrices for each layer */
    for (l = 1; l < nLayers; l++)
        SendMatrix (net [l].wde, lsize [l], lsize [l-1]+1,
            (NM_PTR)mbuf, nodeNo ^ mask, M_WEIGHTS);
}

loopexit:
}

/*----- Now broadcast the complete weight delta matrices to all */
/* nodes, starting with node 0, the root of the tree we used */
/* above. */
stepstowait = maxdiff (nodeNo, 0); /* returns -1 if equal */

```

90/06/06

12:16:18

netnode.c

```

bf [2] = error;
bf [3] = 1;
nwrite ((char *)bf, 4*sizeof(float), host, M_UPDATE, &flag);

for (l = 1; l < nLayers; l++)
{
    mask = (l << 1);
    msgnode = nodeno ^ mask;
    if (l > stepstowait)
    {
        /* accumulate error values from lower level nodes */
        type = M_ERRORS2;
        bf [0] = error;
        nwrite ((char *)bf, sizeof (float), msgnode, type, &flag);
        for (l = 1; l < nLayers; l++)
            SendMatrix (net [l].wdel, lSize[l], lSize [l-1]+1,
                        (MM_PTR)mbuf, nodeno ^ mask, M_WEIGHTS);
    }
    else if (l == stepstowait)
    {
        /* accumulate error values from lower level nodes */
        type = M_ERRORS2;
        nread ((char *)bf, sizeof (float), &msgnode, &type, &flag);
        error = bf [0];
        for (l = 1; l < nLayers; l++)
            GetMatrix (net [l].wdel, lSize[l], lSize [l-1]+1,
                        (MM_PTR)mbuf, nodeno ^ mask, M_WEIGHTS);
    }
}

/* Now adjust weights using entire wdel matrices we just
   received from the broadcast */
for (l = 1; l < nLayers; l++)
    for (in = 0; in < lSize [l]; in++)
        for (j = 0; j <= lSize [l-1]; j++)
            net [l].W [in][j] += net [l].wdel [in][j];

/*----- sum squared Error -----*/
if (MonitorError || (q & nReportErrors == 0))
{
    /* Average error over all patterns */
    error /= (nPatters * net [nLayers-1].lsize);

    if (nodeno == 0)
    {
        bf [0] = (float)q;
        bf [1] = (float)nIterations;
        bf [2] = error;
        bf [3] = 0;
        nwrite ((char *)bf, 4*sizeof(float), host, M_UPDATE, &flag);
    }
    MonitorError = 0;
}

/* Terminate when error satisfactory */
if (error < ErrorLevel)
    break;
}

/*----- end of iteration loop -----*/
if (nodeno == 0)
{
    bf [0] = (float)q;
    bf [1] = (float)nIterations;
}

```

pipe.h

```
/* pipe.h
 */
/* define DEBUG */
/* define QUICKF */
/* define IBM */

#define ESC 27
#define ERRORLEVEL 0.02
#define ITEMS 8
#define MAXLAYERS 5 /* Maximum number of neurons per node */
#define MAXNEURONS 10
#define FALSE 0
#define TRUE 1
#define MAXDATA 12000
#define MAXBITLEN ((MAXDATA+4)*sizeof(float))
#define MAXCUDDIM 6
#define MAX_PTR_MSG 16
#define MS_PER_TICK 0.166
#define LBUFFLEN 512
#define SMALLBUFFLEN 10
#define TEL_SIZE 300

/* message types for inter-processor communication */
#define M_INITIAL 1
#define M_PATTERN 2
#define M_WEIGHTS 3
#define M_UPDATE 4
#define M_OUTPUT 5
#define M_ERRORS 6
#define M_ERRORS2 7
#define M_TIME 8
#define M_UPDATE 9
#define M_SYNCH 10
#define M_TELEGRAM 100

/* typedefs and prototypes for dynamic storage of arrays */
typedef unsigned int UINT;
typedef float *PFLOAT;
typedef PFLOAT VECTOR;
typedef PFLOAT MATRIX;

typedef struct neuron {
    int lNum;
    int nNum;
    VECTOR w;
    VECTOR wde;
} NEURON, *NEURON_PTR;

typedef struct layerrec {
    int lSize;
    int wSize;
    VECTOR del;
    MATRIX wde;
    MATRIX w;
    VECTOR o;
} LAYER;

typedef struct patmsg {
    int p_npatc;
    int p_snum;
    float p_data[MAXDATA];
} PATMSG, *PAT_PTR;

typedef struct vectormsg {
    int v_id;
    int v_cols;
    float v_data[MAXDATA];
} VECTORMSG, *VM_PTR;
```

90/06/06
12/17/02

pipework.c

```
#define PRINTCOND (nodeno <= 3)
/* pipework.c
```

```
Common routines to both the host and node programs. Contains
routines for allocating vectors and matrices, and for sending
vectors and matrices as messages.
```

```
/*
 * stdio.h
 */
#include <stdio.h>
```

```
Copyright (c) Borland International 1987,1988
All Rights Reserved.
```

```
/*
 * pipe.h
 */
#define NULL 0
```

```
#include <string.h>
```

```
#include <ctype.h>
```

```
#include <math.h>
```

```
#include <stdlib.h>
```

```
#include <conio.h>
```

```
#endif
```

```
#include "string.h"
```

```
#include "pipe.h"
```

```
/*
 *----- Vector and Matrix message passing routines -----
 */


```

```
#ifdef IBM
int SendVector(VECTOR vector, int nCols, VM_PTR mbuf, int dest, int msg_id)
```

```
#else
int SendVector(vector, nCols, mbuf, dest, msg_id)
```

```
VECTOR vector;
```

```
int nCols;
```

```
VM_PTR mbuf;
```

```
int dest;
```

```
int msg_id;
#endif
```

```
{
register int i;
```

```
int buflen, flag;
```

```
long result;
```

```
mbuf->v_cols = nCols;
```

```
mbuf->v_id = msg_id;
```

```
#ifdef HOST
swap (&mbuf->v_cols, sizeof (int));
swap (&mbuf->v_id, sizeof (int));
#endif
```

```
for (i = 0; i < nCols; i++)
{
    mbuf->v_data[i] = vector [i];
}
```

```
#ifdef HOST
swap (& (mbuf->v_data[i]), sizeof (float));
#endif
```

```
buflen = 2 * sizeof(int) + nCols * sizeof (float);
```

```
#ifdef HOST
result = nwrite (chan, (char *)mbuf, buflen, dest, msg_id);
return (result);
#endif
```

```
nwrite ((char *)mbuf, buflen, dest, msg_id, &flag);
return (flag);

```

```
#endif
}
```

```
#ifndef IBM
int GetVector(VECTOR vector, int nCols, VM_PTR mbuf, int from, int msg_id)
#else
int GetVector(vector, nCols, mbuf, from, msg_id)
#endif
VECTOR vector;
int nCols;
```

```
VM_PTR mbuf;
```

```
int from;
int msg_id;
#endif
```

```
{
register int i;
```

```
int buflen, type, flag;
```

```
long result;
```

```
buflen = 2 * sizeof(int) + nCols * sizeof (float);
```

```
msgnode = from;
type = msg_id;
#endif
```

```
#ifdef HOST
result = nread (chan, (char *)mbuf, buflen, msgnode, &type);
#endif
```

```
if (type != msgnode)
    nread ((char *)mbuf, buflen, msgnode, &type, &flag);
#endif
```

```
buflen = 2 * sizeof(int) + nCols * sizeof (float);
```

```
msgnode = from;
type = msg_id;
#endif
```

```
#ifdef HOST
result = nread (chan, (char *)mbuf, buflen, msgnode, &type);
#endif
```

```
buflen = 2 * sizeof(int) + nCols * sizeof (float);
```

```
msgnode = from;
type = msg_id;
#endif
```

```
#ifdef HOST
result = swap (& (mbuf->v_data[i]), sizeof (float));
#endif
```

```
vector [i] = mbuf->v_data[i];
#endif
}
```

```
#ifdef HOST
return (result);
#endif
}
#endif

```

```
#ifdef IBM
int SendMatrix(MATRIX matrix, int rows, int cols,
VM_PTR mbuf,
int dest, int msg_id)
#else
int SendMatrix(matrix, rows, cols, mbuf, dest, msg_id)
#endif
MATRIX matrix;
int rows;
int cols;
VM_PTR mbuf;
int dest;
int msg_id;
#endif
```

```
{
register int i;
```

```
int buflen, flag;
```

```
long result;
```

```
mbuf->m_rows = nRows;
```

```
mbuf->m_cols = nCols;
```

```
mbuf->m_id = msg_id;
```

```
#ifdef HOST
swap (&mbuf->m_rows, sizeof (int));
swap (&mbuf->m_cols, sizeof (int));
#endif
```

```
#ifdef HOST
swap (& (mbuf->m_data[i]), sizeof (float));
#endif
```

```
register int i, j;
```

```
int buflen, flag;
```

```
long result;
```

```
#ifdef IBM
int GetMatrix(MATRIX matrix, int rows, int cols,
VM_PTR mbuf,
int from, int msg_id)
#else
int GetMatrix(matrix, rows, cols, mbuf, from, msg_id)
#endif
MATRIX matrix;
int rows;
int cols;
VM_PTR mbuf;
int from;
int msg_id;
#endif
```

```

swap (&mbuf->m_id, sizeof (int));
endif
for (i = 0; i < nRows; i++)
{
    for (j = 0; j < nCols; j++)
    {
        mbuf->m_data[i*nCols+j] = matrix [i][j];
    }
}
#endif HOST
swap (&(mbuf->m_data[1*nCols+1]), sizeof (float));
endif
buflen = 3 * sizeof (int) + nRows * nCols * sizeof (float);
}

one dimensional dynamic vector [cols]
*/
#ifndef IBM
void VectorAllocate (VECTOR *vector, int nCols)
{
    void VectorAllocate (vector, nCols);
    VECTOR *vector;
    int ncols;
}
endif
if (*vector == (VECTOR) calloc (nCols, sizeof (float))) == NULL)
{
}
#endif
/* Allocate space for columns (float cells) for
dynamic two dimensional matrix [rows] [cols]
*/
#ifndef IBM
void AllocateCols (PFLOAT matrix[], int nRows, int nCols)
{
    void AllocateCols (matrix, nRows, nCols);
    PFLOAT matrix[];
    int nrows;
    int ncols;
}
endif
int i;
for (i = 0; i < nRows; i++)
{
    VectorAllocate (matrix[i], ncols);
}
#endif
/* Allocate space for a two dimensional dynamic matrix [rows] [cols]
*/
#ifndef IBM
void MatrixAllocate (MATRIX *pmatrix, int nRows, int nCols)
{
    void MatrixAllocate (pmatrix, nRows, nCols);
    MATRIX *pmatrix;
    int nrows;
    int ncols;
}
endif
if (*pmatrix == (MATRIX) calloc (nRows, sizeof (FLOAT)) ) == NULL)
{
}
AllocateCols (*pmatrix, nrows, ncols);
#endif
/* free space for two dimensional dynamic array */
#ifndef IBM
void MatrixFree (MATRIX matrix, int nRows)
{
    void MatrixFree (matrix, nRows);
    MATRIX matrix;
    int nrows;
}
endif
int i;
for (i = 0; i < nRows; i++)
{
    for (j = 0; j < nCols; j++)
    {
        register int i, j;
        int buflen, msgnode, type, flag;
        long result;
        buflen = 3 * sizeof (int) + nRows * nCols * sizeof (float);
        msgnode = from;
        type = msg_id;
    }
}
#endif HOST
result = nread (chan, (char *)mbuf, buflen, msgnode, type, flag);
if (i = 0; i < nRows; i++)
{
    for (j = 0; j < nCols; j++)
    {
        if (*mbuf->m_data[1*nCols+j]) , sizeof (float));
    }
}
#endif
swap (&(mbuf->m_data[1*nCols+1]), sizeof (float));
endif
matrix [i][j] = mbuf->m_data[i*nCols+j];
endif
return (result);
else
    return (flag);
}
#endif
----- Array storage allocation routines -----
/* Allocate space for vector of float cells for
-----
```

90/06/06
12:17:02

pipework.c

```

free(matrix[1]);
free(matrix);

LAYER *LayerAllocate(LAYER *layer, int layersize, int wtsize)
{
    ifdef IBM
    void LayerAllocate(LAYER *layer, layersize, wtsize)
    else
    void LayerAllocate(layer, layersize, wtsize)
    LAYER *layer;
    int layersize;
    int wtsize;
    endif
    {
        layer->size = layersize;
        layer->wSize = wtsize;
        /* w and wdel must contain room for the node bias */
        MatrixAllocate (&(layer->w), layersize, wtsize+1);
        MatrixAllocate (&(layer->wdel), layersize, wtsize+1);
        VectorAllocate (&(layer->o), layersize);
        VectorAllocate (&(layer->odel), layersize);
    }

    ifdef IBM
    void LayerFree(LAYER *layer)
    else
    void LayerFree(layer)
    LAYER *layer;
    endif
    {
        MatrixFree (layer->w, layer->size);
        MatrixFree (layer->odel, layer->size);
        free (layer->o);
        free (layer->odel);
    }

    ifdef IBM
    float f(float x)
    else
    float f(x)
    float x;
    endif
    {
        ifdef QUICKF
        static int doneInit = FALSE;
        static VECTOR v;
        static int vsize = 1000;
        static float lo = -5.0;
        static float hi = 5.0;
        static float step;
        register int i;
        float y, y2, d;
        if (!doneInit)
        {
            VectorAllocate (&v, vsize);
            step = (hi - lo) / vsize;
            for (i = 0, y = lo; i < vsize; i++, y+=step)
                v[i] = 1.0 / (1.0 + exp {-y});
            doneInit = TRUE;
        }
        if (x < lo || x >= hi)
            return (1.0 / (1.0 + exp {-x}));
        else
        {
            i = (x - lo) / step;
            y = (x - lo + i * step) / step;
            return (v[i] + y * (v[i+1] - v[i]));
        }
    }
}

void LandN (lSize, num, L, N)
{
    int lSize[],
        int num,
        int *L,
        int *N;
    for (i = 1, s = 0; s <= num; i++)
        s += lSize[i];
}

```

3

```
*L = l-1;
*N = num - (s - 1Size [l-1]);
}

#ifndef NODE
void send_synch(node)
int node;
{
    int l, flag;

    l = 1;
    nwrite (&l, sizeof (int), node, M_SYNCH, flag);
}
#endif

void get_synch(node)
int node;
{
    int from, type, flag, l;

    type = M_SYNCH;
    from = node;
}
#ifndef NODE
    nread (&l, sizeof (int), &from, &type, &flag);
#endif
#ifndef HOST
    printf ("Waiting for synchronization from %d...", from);
    nread (chan, &i, sizeof (int), &from, type);
    printf ("Got it..\n");
#endif
}
#endif

#ifndef NODE
int dprintf ()
{
    int flag;

    if (PRINTCOND)
        nwrite ((char *)dbuf, TEL_SIZE * sizeof (char), host, M_TELEGRAM, &flag);
}
#endif
```


pipehost.c

```

    /* learning rate */
    /* momentum factor */
    /* number of layers in network */

    /* get sizes of layers */
    nNeurons = 0;
    for (l = 0; l < nlayers; l++)
    {
        fscanf(fpinp, "%d", &lSize [l]);
        if (l != 0) nNeurons += lSize [l];
    }

    /*-----Allocate dynamic storage for all data -----*/
    LayerAllocate(&net [1], lSize [1], {l=0} ? 0:lSize [l-1]);
    MatrixAllocate(&target,
    nPatterns,
    lSize [nlayers-1]);
    MatrixAllocate(&input,
    nPatterns,
    lSize [0]);
    VectorAllocate(&PatternID, nPatterns);

    /*----- Read the initial weight matrices: -----*/
    if ((fpWeights = fopen(szWeights, "r")) == NULL)
    {
        fprintf(stderr, "can't open file %s\n",
            progname, szWeights);
        exit(1);
    }

    /* read weights */
    for (l = 1; l < nlayers; l++)
        for (n = 0; n < net [l].lSize; n++)
            for (i = 0; i < net [l-1].lSize; i++)
            {
                fscanf(fpWeights, "%f",
                    &net [l].w[n][i]);
                net [l].wdel[n][i] = 0.0;
            }
    fclose(fpWeights);

    /*----- Read in all patterns to be learned -----*/
    for (p = 0; p < nPatterns; p++)
    {
        for (l = 0; l < net [0].lSize; l++)
            if (fscanf(fpinp, "%f",
                &input [p] [l]) != 1)
                goto ALLPATTERNSREAD;

        /* read in target outputs for input patterns read */
        for (l = 0; l < net [nlayers-1].lSize; l++)
            fscanf(fpinp, "%f",
                &target [p] [l]);
    }

    fscanf(fpinp, "%f", &PatternID[p]);
}

/* All input data has been read. Now start sending information to the nodes */
cubedim = MAXCUBDIM + 1;
while ((cubedim < 0) || (cubedim > MAXCUBDIM)) {
    printf("Enter Cube Dimension (0 - %d : ", MAXCUBDIM);
    scanf("%d", &cubedim);
}

cubsize = power (2, cubedim); /* number of processors */
lbuf = (char *)malloc(LBUFSIZE);
printf ("Attempting to open cube: Try ");
i = 0;
do {
    printf ("%d.", i++);
    while ((chan = fopen(cubedim)) < 0);
    printf ("Got channel %d\n", chan);
}

/* Load the node program */
printf ("Loading node program.\n");
nload(chan, "node.x", -cubesize, nflags, lbuf, LBUFSIZE);
free(lbuf);
printf ("Node program loaded to %d processors.\n", cubesize);

fprintf(stderr, "Setting up initial message to nodes.\n");
((IM_PTR)lbuf)->l_layers = nLayers;
((IM_PTR)lbuf)->l_iters = nIterations;
((IM_PTR)lbuf)->l_pats = nPatterns;
((IM_PTR)lbuf)->l_seed = 1L;
((IM_PTR)lbuf)->l_eta = eta;
((IM_PTR)lbuf)->l_alpha = alpha;
((IM_PTR)lbuf)->l_epsilon = ErrorLevel;
printf ("Set up initial message. Now swapping.\n");
swap (&((IM_PTR)lbuf)->l_layers, sizeof(int));
swap (&((IM_PTR)lbuf)->l_iters, sizeof(int));
swap (&((IM_PTR)lbuf)->l_pats, sizeof(int));
swap (&((IM_PTR)lbuf)->l_seed, sizeof(long));
swap (&((IM_PTR)lbuf)->l_eta, sizeof(float));
swap (&((IM_PTR)lbuf)->l_alpha, sizeof(float));
swap (&((IM_PTR)lbuf)->l_epsilon, sizeof(float));

for (l = 0; l < MAXLAYERS; l++)
{
    ((IM_PTR)lbuf)->l_size [l] = lSize [l];
    swap (&((IM_PTR)lbuf)->l_size[l], sizeof(int));
}

near_node = (float)nNeurons / (cubesize - 1) /* -1 for splicer */;
near_num = 0;
printf ("Now sending initial message to nodes. Total neurons = %d\n",
    nNeurons);

for (l = 0; l < cubesize; l++)
{
    /* Compute left and right neighbors of node l */
    L = bintogray ((l==0) ? cubesize-1 : l-1);
    R = bintogray -(l+1) & cubesize;
    ((IM_PTR)lbuf)->l_nbr = L;
    ((IM_PTR)lbuf)->l_rnbr = R;
    /* Now calculate which neurons node l will process */
    if (l == 0)

```

pipehost.c

```

startN = (int) neur_num;
endN = (int) (neur_num + neur_node - 1);
if (1 == cubeSize)
    endN = nNeurons - 1;
    ((IM_PTR)ibuf)->1_startN = startN;
    ((IM_PTR)ibuf)->1_endN = endN;
    neur_num += neur_node;
}
else
{
    startN = -1;
    endN = -1;
}

printf ("node %d (%d) gets neurons %d to %d.\n", i,
    ((IM_PTR)ibuf)->1_Lnbr,
    ((IM_PTR)ibuf)->1_Rnbr,
    ((IM_PTR)ibuf)->1_startN,
    ((IM_PTR)ibuf)->1_endN);
swap (&((IM_PTR)ibuf)->1_startN, sizeof(int));
swap (&((IM_PTR)ibuf)->1_endN, sizeof(int));
swap (&((IM_PTR)ibuf)->1_endN, sizeof(int));
swap (&((IM_PTR)ibuf)->1_Rnbr, sizeof(int));
swap (&((IM_PTR)ibuf)->1_Rnbr, sizeof(int));

/* Send initial message to node */
type = M_INITIAL;
nwrite (chan, (char *)ibuf, sizeof(INITIALMSG), bintogray (1), type);
if (1 != 0) /* If not the splicer, send the weights */
for (j = startN; j < endN; j++)
{
    printf ("Finding l and n for %d\n", j);
    LandN (lsize, j, &l, &n);
    printf ("Sending weights for %d, %d to %d.", l, n, bintogray (1));
    sendVector (net [1].W [l], lsize [l-1]+1, (VM_PTR)ibuf,
        bintogray (1), M_WEIGHTS);
    for (d1=0; d1 < 100; d1++)
        for (d2=0; d2 < 1000; d2++) /* nothing */
    printf ("Done..\n");
}

else /* send the input and target output matrices to the splicer */
{
    SendMatrix (input, nPatterns, lsize [0], (MM_PTR)ibuf,
        1, M_PATTERN);
    SendMatrix (target, nPatterns, lsize [nLayers-1], (MM_PTR)ibuf,
        1, M_OUTPUT);
}

printf ("Waiting for synchronization from %d..\n", bintogray (1));
Get_Synch (bintogray (1));
printf ("Got it.\n");

printf ("Awaiting results..\n");
/* Now we await the results */
tot_time = 0;
time_cnt = 0;
j = 0;
i = 0;
do
{
    char c;
    long l;
    type = -1;
    while (intest (chan, &i, &type) < 0)
}

```

98/06/06
12:16:42

pipehost.c

```
int nrows;
int ncols;
#endif
{
    int i, j;
    for (i = 0; i < nrows; i++)
        for (j = 0; j < ncols; j++)
            printf ("%5.4f%c", m[i][j], (j==ncols-1) ? '\n' : ' ');
    #ifdef IBM
    if (i & 25 == 24) getch();
    #endif
}
return (1);
}
```

900606
12:16:50

pipenode.c

```
* pipenode.c
Jeffrey Dean
ID 1321294
Cscl 5299
Term Project

Pipelined back-propagation neural network

#define NODE

#ifndef IBM
#include <stdlib.h>
#endif
#include <math.h>
#include <ctype.h>
#include <string.h>
#include "pipe.h"

/* define storage for net layers */
/* Arrays for inputs, outputs, deltas, weights & targets */
nt lsize[MAXLAYERS]; /* number of patterns desired */
nt nPatterns, /* Number of patterns in this message */
nNeurons, /* Number of iterations desired */
iterations, /* starting neuron number for node */
startN, /* ending neuron number for node */
endN, /* number of neurons for this node */
numN, /* number of layers in network */
nLayers, /* maximum layer size */
maxLSize;
ECTOR target;
ECTOR in;
ECTOR o;
ECTOR del[MAXLAYERS];
EURON N[MAXNEURONS];
/* The following three matrices are only used on the splicer node */
ATRIX targetM;
ATRIX inputM;
ATRIX deltaM;
nt msg_stack; /* number of messages that need to be added to pipeline */
nt patWidth;
nt patMsg_size;
num_msgs;
work_nodes;
pat_par_msg;
max_pats;
har;
done;

nt epoch;
patr;
loat eta = 0.15;
alpha = 0.075;
loat error;
loat lastError;
ong ranseed = 12357341;

/* Cube communication variables */
nt nodeno, proc, host, cubedim, flag;
nt har;
nt b [SMALLBUFSIZE];
loat left, right;
nt nttype;

* include "pipework.c"
void init_pointers();
void rotate_level();
void unroll();
void process_neuron();
void update_weights();
void setup_pat_msg();
void send_pat_msg();
get_pat_msg();
void send_update_msg();
void get_update_msg();
void print_vec (s, v, sz)
char *s;
float v [];
int sz;
{
    int i;
    dbuf[0] = '\0';
    i = strlen(dbuf);
    sprintf(dbuf, "%s Len = %d %d %4.2f %4.2f %4.2f", s, i, sz, v[0], v[1], v[2]);
    dprintf();
    sprintf (dbufz, "%s (%d): %s, %s, %s");
    for (i = 0; i < sz; i++)
        sprintf (dbufz+strlen (dbuf), ", %4.2f", v[i]);
    dprintf();
}

void print_buffer (pNum)
int pNum;
{
    int i;
    init_pointers ((PAT_PTR)dbuf, pNum);
    print_vec ("Input", ln, lsize [0]);
    print_vec ("Target", target, lsize [nLayers-1]);
    print_vec ("Msg", ((PAT_PTR)dbuf) ->p_data[pNum*pat_msg_size], pat_msg_size);
}

void main()
{
    int reportErrors = 1; /* error reporting frequency */
    float ErrorLevel = ERRORLEVEL; /* satisfactory error level */
    char MonitorError = 0; /* true when monitor error display */
    int i;
    int n, p, Q;
    /* neuron index
     * Index pattern number
     * Index iteration number
     */
    int srime;
    whoami (fnodeno, sproc, ghost, &cubedim);
    work_nodes = {1 << cubedim} - 1;
    /* allocate message buffer */
    mbuf = (char *)malloc(MAXBUFSIZE);
    type = M_INITIAL;
```

20/06/06
12:16:50

pipenode.c

```

read ( (char *)mbuf, sizeof (INITIALMSG), &host, &type, &flag );
if (nodeNo != 0)
{
    sprintf (dbuf, "Waiting for weight vectors from %d neurons.", numN);
    dprintf ();
    /****** Pipeline node section ***** */
    /* Allocate weight vectors for each neuron and get weights from host */
    for (n = 0; n < numN; n++)
    {
        int lev, neur;
        sprintf (dbuf, "Trying to get L and N for %d", n);
        dprintf ();
        levN (n + startN, elev, neur);
        N[n].lNum = neur;
        N[n].lSize [lev] = lev;
        sprintf (dbuf, "Waiting for %d (%d,%d), size = %d",
                n, lev, neur, lsize [lev-1]+1);
        dprintf ();
        VectorAllocate (N[n].w, lsize [lev-1]+1);
        VectorAllocate (N[n].wdel, lsize [lev-1]; 1+1);
        for (i = 0; i <= lsize [lev-1]; 1++)
        {
            N[n].wdel [i] = 0;
            GetVector (N[n].w, lsize [lev-1]+1, (WM_PTR)mbuf, host, M_WEIGHTS);
        }
        send_synch (host);
        /* Send synchronizing message to host */
        sprintf (dbuf, "Got weights for %d neurons, npats=%d (%d, %d",
                numN, npatterns, left, right);
        dprintf ();
    }
    /* Main loop for pipeline nodes */
    while (1)
    {
        mtype = get_message ();
        switch (mtype)
        {
            case M_UPDATE:
                get_update_msg ();
                sprintf (dbuf, "Got update message %d of %d epochs done",
                        (UM_PTR)mbuf->u_iter, ((UM_PTR)mbuf)->u_maxiter);
                dprintf ();
                send_update_msg (0, 0);
                update_weights ();
                if (((UM_PTR)mbuf)->u_iter >= ((UM_PTR)mbuf)->u_maxiter)
                    done = TRUE;
                break;
            case M_PATTERN:
                get_pat_msg ();
                for (p = 0; p < ((PAT_PTR)mbuf)->p_npat; p++)
                    init_pointers ((PAT_PTR)mbuf, p);
                for (n = 0; n < numN; n++)
                {
                    process_neuron (n);
                    if (N[n].lNum == lsize [N[n].lNum] - 1)
                        rotate_level (N[n].lNum, p);
                }
                send_pat_msg ();
                break;
            default:
                sprintf (dbuf, "Unexpected message of type %d", mtype);
                dprintf ();
                done = TRUE;
                break;
        }
    }
}
/* get starting time */
stime = nttime();
/* get starting time */
done = FALSE;
epoch = 0;
done = FALSE;

```

pipenode.c

```

/*
 * sprintf (dbuf, "Got back and rolled up message for id to %d",
 *          ((PAT_PTR)mbuf)->p_snum,
 *          ((PAT_PTR)mbuf)->p_nptrs);
 *
 * msg_stack++; /* room for one more message in the pipeline */
 * break;
 * default:
 *   sprintf (dbuf, "Unexpected message of type %d", mtype);
 *   dprintf();
 *   done = TRUE;
 *   break;
 * }
 * send_synch (host); /* Synchronize with host */
 /* End of while */
 **** End splicer section ****
 */

/* Initialize delras to 0 */
for (p = 0; p < npatterns; p++)
  for (i = 0; i < delrawidth; i++)
    deltarM[p][i] = 0;

sprintf (dbuf, "Splicer has been initialized. delrawidth = %d, delrawidth");
dprintf();

ptr = 0;
msg_stack = work_nodes; /* set to fill pipeline initially */
lasterror = 100.0;
error = 0;

while (!done)
  while ((msg_stack > 0) && (!done))
    {
      sprintf (dbuf, "Sending patterns %f to %f (epoch %d)", ptr,
              iTr + pat_per_msg, epoch);
      iTr += pat_per_msg;
      ptr += pat_per_msg;
      if (ptr >= nPatterns)
        {
          send_update_msg (lasterror);
          ptr = 0;
          if (epoch++ == nIterations)
            done = TRUE;
        }
      msg_stack--;
    }

if (!done)
{
  mtype = get_message();
  switch (mtype)
  {
    case M_UPDATE:
      get_update_msg();
      sprintf (dbuf, "Got weight update back for epoch %d, epoch");
      dprintf();
      break;
    case M_PATTERN:
      get_pat_msg();
      rollup();
  }
}
else
{
  /* Allocate and get input and expected output matrices */
  MatrixAllocate (inputM, nPatterns, lsize [0]);
  MatrixAllocate (targetM, nPatterns, lsize [nLayers-1]);
  GetMatrix (inputM, nPatterns, lsize [0], (MM_PTR)mbuf, host, M_PATTERN);
  GetMatrix (targetM, nPatterns, lsize [nLayers-1], (MM_PTR)mbuf, host,
             M_OUTPUT);
  send_synch (host); /* Synchronize with host */
  /* Allocate matrix to hold delta vectors */
  MatrixAllocate (deltaM, nPatterns, delrawidth);
  sprintf (dbuf, "Splicer section ****
 */

int maxdiff (n1, n2)
int n1, n2;
{
  int i, j, mask;
  for (i = 0, j = 0; i < 8 * sizeof (int); i++)
    {
      mask = (1 << i);
      if ((n1 & mask) != (n2 & mask))
        j++;
    }
  if (n1 == n2)
    return (-1);
  else
    return (j);
}

void init_pointers (pbuf, pNum)
PAT_PTR pbuf;
int pNum;
{
  /* Init_pointers calculates the offsets into the pattern buffer for
   * the target vector, input and output vectors, and the delta vectors.
   * The data is arranged as follows:
   */
  Offset      Data
  0:          Target vector
  $[nL-1]:    Delta vector 1
  +$[1]:      Input vector
  +maxSize:   Output vector
  ...
  +$[nL-2]:   Delta vector nL-1
  +$[nL-1]:   Input vector
  +maxSize:   Output vector
  [
    int i, temp, t2;
    /* Calculate offset of pth pattern into message buffer */
}

```

```

t2 = pat_msg_size * pNum;

/* Now calculate offsets into pattern data array */
target = &(pbuf->p_data[0+t2]);
temp = lSize [nLayers-1]; /* Target is output vector for last layer */
for (l = 1; l < nLayers; l++)
{
    del[l] = *(pbuf->p_data[t2+l]);
    temp += lSize [l];
}
in = &(pbuf->p_data[temp+t2]);
o = &(pbuf->p_data[temp+maxLSize*t2]);
}

old rotate_level (from_lev, pNum)
{
    int from_lev;
    int pNum;
    int i;

    if (from_lev != nLayers - 1)
    {
        /* Initialize pointers for this pattern */
        init_pointers ((PAT_PTR)mbuf, pNum);

        /* First step: copy output vector to input vector */
        for (i = 0; i < maxLSize; i++)
            in [i] = o [i];

        /* Now zero the delta vector for the level we're coming from */
        for (i = 0; i < lSize [from_lev]; i++)
            del[from_lev][i] = 0;
    }
}

old process_neuron (NN)
{
    int i, l, n;
    float sum, dw;
}

/* Get nNum and lNum into n and l for convenience */
n = NN.nNum;
l = NN.lNum;

**** Step I. Calculate new activation values given the input vector */
sum = NN.W[lSize[l-1]];

for (i = 0; i < lSize [l-1]; i++)
    sum += in [i] * NN.W [i];
o [i] = f (sum);

**** Step II. Calculate deltas */
/* If we're an output layer, we have to first calculate the error from
   the expected output */
if (l == nLayers - 1)
    del [1][n] = (target [n] - o [n]);
/* Now scale the sum of the errors by the derivative of the f function */
del [1][n] *= o [n] * (1.0 - o [n]);
if (l != 1) /* if not the first hidden layer */
{
    del [1][n] *= o [n] * (1.0 - o [n]);
}

/* Now back-propagate using the delta values from the current layer */
for (i = 0; i < lSize [l-1]; i++)
    del [1][i] += NN.W [i] * del [1][n];

**** Step III. Calculate weight changes based on deltas we just calculated */
dw = eta * del [1][n];
NN.W[lSize [l-1]] += dw;

for (i = 0; i < lSize [l-1]; i++)
{
    dw = eta * del [1][i] + in [i];
    /* Input to this layer is output from the previous layer */
    NN.W[i].wdel [i] += dw;
}

}

void unroll (p, pnum)
{
    int p7, /* slot number to put pattern in */
    int pnum7 /* pattern to be unrolled */;

    int i;

    /* First we have to initialize the pointers to the correct locations */
    init_pointers ((PAT_PTR)mbuf, p);
    /* Copy target pattern to target vector */
    for (i = 0; i < lSize [nLayers-1]; i++)
        target [i] = targetM [p+i];
    /* Copy delta vectors. Note that deltas are contiguous in the message
       buffer, so we can just copy it as if it was one vector. */
    for (i = 0; i < delWidth; i++)
        del [1][i] = deltaM [p+i];
    /* Copy input pattern to activation vector and zero rest of activation
       vector and all of the output vector. */
    for (i = 0; i < maxLSize; i++)
    {
        in [i] = (i < lSize [0] ? inputM [p][i] : 0);
        o [i] = 0;
    }
}

void rollup ()
{
    int i, p, npats, sp;
    float temp;
    /* Temporary value to hold error term */

    npats = ((PAT_PTR)mbuf)->p_npats;
    sp = ((PAT_PTR)mbuf)->p_nspats;
    for (p = 0; p < npats; p++)
    {
        init_pointers ((PAT_PTR)mbuf, p);
        for (i = 0; i < delWidth; i++)
            del [1][i] += deltaM [p+i];
        for (i = 0; i < lSize [nLayers-1]; i++)
        {
            temp = target [i] - o [i];
            error += temp * temp;
        }
        if (sp + p == nPatterns - 1) /* If last pattern in the epoch.. */
        {
            lasterror = error / (nPatters * lSize [nLayers-1]);
            sprintf (dbuf, "END of EPOCH %d: Error = %f", epoch, lasterror);
            dprintf ();
            error = 0; /* Initialize error for next epoch */
        }
    }
}

```

98/06/06
12:16:50

pipenode.c

```
result = ntest (aleft, mttype);
        while (result < 0);
        return (mttype);
}

void update_weights()
{
    int l, n;

    for (n = 0; n < numN; n++)
    {
        /* sprintf (dbuf, "Weight delta for %d: ", n+startN); */
        for (l = 0; l <= lsize [N[n].lnum - 1]; l++)
        {
            N[n].w[l] += N[n].wdel[l];
            /* sprintf (dbuf+strlen(dbuf), "%f ", N[n].wdel[l]); */
            N[n].wdel[l] *= 0; /* keep momentum */
        }
        /* dprintf(); */
    }
}

int npats;
int p, npats;
((PAT_PTR)mbuf)->p_snum = (int)start_pat;
for (p = (int)start_pat; p < end_pat; p++)
if (p < nPatterns)
{
    unroll (p, npats);
    npats++;
}

id send_pat_msg()
{
    int msgsize;
    msgsize = pat_msg_size * max_pats * sizeof (float) + 2 * sizeof (int);
    nwrite ((char *)mbuf, msgsize, right, M_PATTERN, &flag);
}

id get_pat_msg()
{
    int msgsize, mttype;
    mttype = M_PATTERN;
    msgsize = pat_msg_size * max_pats * sizeof (float) + 2 * sizeof (int);
    nread ((char *)mbuf, msgsize, aleft, mttype, &flag);

    id send_update_msg()
    {
        if (nodeo == 0) /* If we're the splicer, set up the message */
        {
            ((UM_PTR)mbuf)->u_iter = epoch;
            ((UM_PTR)mbuf)->u_maxiter = nIterations;
            ((UM_PTR)mbuf)->u_error = e;
            ((UM_PTR)mbuf)->u_error = e;
        }
        nwrite ((char *)mbuf, sizeof (UPDATEMSG), right, M_UPDATE, &flag);
    }

    id get_update_msg()
    {
        int mttype;
        mttype = M_UPDATE;
        nread ((char *)mbuf, sizeof (UPDATEMSG), aleft, mttype, &flag);
    }

    t get_message()
    int result, mttype;
    do {
        mttype = -1;
```